
Adversarial Mutual Information for Text Generation

Boyuan Pan^{1*} Yazheng Yang^{2*} Kaizhao Liang³ Bhavya Kailkhura⁴ Zhongming Jin⁵ Xian-Sheng Hua⁵
Deng Cai¹ Bo Li³

Abstract

Recent advances in maximizing mutual information (MI) between the source and target have demonstrated its effectiveness in text generation. However, previous works paid little attention to modeling the backward network of MI (*i.e.* dependency from the target to the source), which is crucial to the tightness of the variational information maximization lower bound. In this paper, we propose *Adversarial Mutual Information* (AMI): a text generation framework which is formed as a novel saddle point (min-max) optimization aiming to identify joint interactions between the source and target. Within this framework, the forward and backward networks are able to iteratively promote or demote each other’s generated instances by comparing the real and synthetic data distributions. We also develop a latent noise sampling strategy that leverages random variations at the high-level semantic space to enhance the long term dependency in the generation process. Extensive experiments based on different text generation tasks demonstrate that the proposed AMI framework can significantly outperform several strong baselines, and we also show that AMI has potential to lead to a tighter lower bound of maximum mutual information for the variational information maximization problem.

1. Introduction

Generating diverse and meaningful text is one of the coveted goals in machine learning research. Most sequence transductive models for text generation can be efficiently trained by maximum likelihood estimation (MLE) and have

*Equal contribution ¹State Key Lab of CAD&CG, Zhejiang University ²College of Computer Science, Zhejiang University ³University of Illinois Urbana-Champaign ⁴Lawrence Livermore National Laboratory ⁵Alibaba Group. Correspondence to: Boyuan Pan <panby@zju.edu.cn>, Bo Li <lbo@illinois.edu>.

demonstrated dominant performance in various tasks, such as dialog generation (Serban et al., 2016; Park et al., 2018), machine translation (Bahdanau et al., 2014; Vaswani et al., 2017) and document summarization (See et al., 2017). However, the MLE-based training schemes have been shown to produce safe but dull texts which are ascribed to the relative frequency of generic phrases in the dataset such as “*I don’t know*” or “*what do you do*” (Serban et al., 2016).

Recently, Li et al. (2016a) proposed to use maximum mutual information (MMI) in the dialog generation as the objective function. They showed that modeling mutual information between the source and target will significantly decrease the chance of generating bland sentences and improve the informativeness of the generated responses, which has inspired many important works on MI-prompting text generation (Li et al., 2016b; Zhang et al., 2018b; Ye et al., 2019). In practice, directly maximizing the mutual information is intractable, so it is often converted to a lower bound called the variational information maximization (VIM) problem by defining an auxiliary backward network to approximate the posterior of the dependency from the target to the source (Chen et al., 2016). Specifically, the likelihood of generating the real source text by the backward network, whose input is the synthetic target text produced by the forward network, will be used as a reward to optimize the forward network, tailoring it to generate sentences with higher rewards.

However, most of them optimize the backward network by maximizing the likelihood stated above. Without the guarantee of the quality of the synthetic target text produced by the forward network, this procedure encourages the backward model to generate the real source text from many uninformative sentences and is not likely to make the backward network approach the true posterior, thus affects the tightness of the variational lower bound. On the other hand, since the likelihood of the backward network is a reward for optimizing the forward network, a biased backward network may provide unreliable reward scores, which will mislead the forward network to still generate bland text.

In this paper, we present a simple yet effective framework based on the maximum mutual information objective that encourages it to train the forward network and the backward network adversarially. We propose *Adversarial Mu-*

tual Information (AMI), a new text generative framework that addresses the above issues by solving an information-regularized minimax optimization problem. Instead of maximizing the variational lower bound for both the forward network and the backward network, we maximize the objective when training the forward network while minimizing the objective when training the backward network. In addition to the original objective, we add the expectation of the negative likelihood of the backward model whose sources and targets are sampled from the real data distribution. To this end, we encourage the backward network to generate real source text only when its input is in the real target distribution. This scheme remedies the problem that the backward model gives high rewards when its input is the bland text or the text that looks human-generated but has little related content to the real data. To stabilize training, we prove that our objective function can be transformed to minimizing the Wasserstein Distance (Arjovsky et al., 2017), which is known to be differentiable almost everywhere under mild assumptions. Moreover, we present a latent noise sampling strategy that adds sampled noise to the high-level latent space of the model in order to affect the long-term dependency of the sequence, hence enables the model to generate more diverse and informative text.

Technical Contributions In this paper, we take the first step towards generating diverse and informative text by optimizing adversarial mutual information. We make contributions on both the theoretical and empirical fronts.

- We propose *Adversarial Mutual Information* (AMI), a novel text generation framework that adversarially optimizes the mutual information maximization objective function. We provide a theoretical analysis on how our objective can be transformed into the Wasserstein distance minimization problem under certain modifications.
- To encourage generating diverse and meaningful text, we develop a latent noise sampling method that adds random variations to the latent semantic space to enhance the long term dependency for the generation process.
- We conduct extensive experiments on the tasks of dialog generation and machine translation, which demonstrate that the proposed AMI framework outperforms several strong baselines significantly. We show that our method has the potential to lead to a tighter lower bound of maximum mutual information.

2. Background

2.1. Sequence Transduction Model

The sequence-to-sequence models (Cho et al., 2014; Sutskever et al., 2014) and the Transformer-based mod-

els (Vaswani et al., 2017) have been firmly established as the state-of-the-art approaches in text generation problems such as dialog generation or machine translation (Wu et al., 2019; Pan et al., 2018; Akoury et al., 2019). In general, most of these sequence transduction models hold an encoder-decoder architecture (Bahdanau et al., 2014), where the encoder maps an input sequence of symbol representations $S = \{x_1, x_2, \dots, x_n\}$ to a sequence of continuous representations $z = \{z_1, z_2, \dots, z_n\}$. Given z , the decoder then generates an output sequence $T = \{y_1, y_2, \dots, y_m\}$ of tokens one element at a time. At each time step the model is auto-regressive, consuming the previously generated tokens as additional input when generating the next.

For the sequence-to-sequence models, the encoder and the decoder are usually based on recurrent or convolutional neural networks, and the Transformer builds them using stacked self-attention and point-wise fully connected layers. Given the source S and the target T , the objective of a sequence transduction model can be simply formulated as:

$$\theta^* = \arg \max_{\theta} P_{\theta}(T|S) \quad (1)$$

where θ is the parameters of the sequence transduction model.

2.2. Maximum Mutual Information

Maximum mutual information for text generation is increasingly being studied (Li et al., 2016a; Zhang et al., 2018b; Ye et al., 2019). Compared to the maximum likelihood estimation (MLE) objective, maximizing the mutual information $I(S, T)$ between the source and target encourages the model to generate sequences that are more specific to the source. Formally, we have

$$I(S, T) = \mathbb{E}_{P(T,S)} \left[\log \frac{P(S, T)}{P(S)P(T)} \right]. \quad (2)$$

However, directly optimizing the mutual information is intractable. To provide a principled approach to maximizing MI, a Variational Information Maximization (Barber & Agakov, 2003; Chen et al., 2016; Zhang et al., 2018b) lower bound is adopted:

$$\begin{aligned} I(S, T) &= H(S) - H(S|T) \\ &= H(S) + \mathbb{E}_{P(T,S)} [\log P(S|T)] \\ &= H(S) + \mathbb{E}_{P(T)} [D_{KL}(P(S|T) || Q_{\phi}(S|T))] \\ &\quad + \mathbb{E}_{P(T,S)} [\log Q_{\phi}(S|T)] \\ &\geq H(S) + \mathbb{E}_{P(S)} \mathbb{E}_{P_{\theta}(T|S)} [\log Q_{\phi}(S|T)] \end{aligned} \quad (3)$$

where $H(\cdot)$ denotes the entropy, and $D_{KL}(\cdot, \cdot)$ denotes the KL divergence between two distributions. $Q_{\phi}(S|T)$ is a *backward network* that approximates the unknown $P(S|T)$, and usually shares the same architecture with the forward

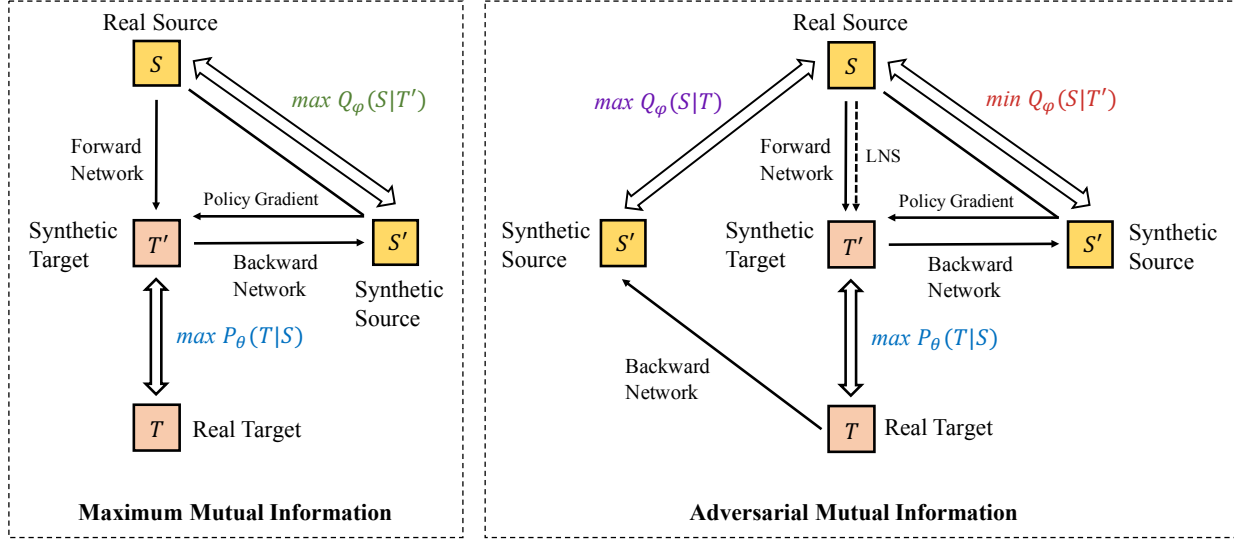


Figure 1. The overview of frameworks for the Maximum Mutual Information (left part) and the proposed Adversarial Mutual Information (right part). Different colors of text denotes different objectives.

network $P_\theta(T|S)$. Note that the entropy term associated with the training data is ignored as it does not involve the parameters we are trying to optimize.

Given the source text S , the objective of MMI is equivalent to maximizing:

$$\max_{\theta, \phi} \mathbb{E}_{T' \sim P_\theta(T'|S)} [\log Q_\phi(S|T')] \quad (4)$$

where θ and ϕ are the parameters of the forward network and the backward network respectively. We denote the synthetic target text generated by the forward network as T' in order to distinguish it from the real data.

3. Adversarial Mutual Information

In this section, as shown in the Figure 1, we present the drawbacks of directly optimizing the MMI problem and introduce how our proposed Adversarial Mutual Information (AMI) solves these issues by adversarially optimizing its modules.

Intuitively, the objective in the equation (4) for the forward network P_θ can be viewed as to use the log-likelihood of the backward model Q_ϕ , whose input is the synthetic target text produced by P_θ , to reward or punish P_θ . This encourages the forward network to generate text that has a stronger connection with its input. For the backward network Q_ϕ , however, maximizing the log-likelihood of Q_ϕ when its input is the synthetic text is not always optimal, especially when this synthetic text is uninformative or has little intersection in content with the source text. Although this scheme is maximizing the mutual information in the form, it optimizes Q_ϕ towards a biased direction and a low-quality

Q_ϕ may bring unreliable reward signal. In addition, the lower bound shown in the equation (3) becomes tight as the approximated distribution Q_ϕ approaches the true posterior distribution: $\mathbb{E} [D_{KL}(P(\cdot)|Q_\phi(\cdot))] \rightarrow 0$, which is not likely to be achieved when many negative samples are involved.

Inspired by the delicate but straightforward idea of the generative adversarial networks (Goodfellow et al., 2014), we set all the synthetic target text as the negative samples and force the backward model to generate correct source text only if its input is the real target text. Formally, we propose to solve the following minimax game:

$$\min_{\phi} \max_{\theta} \underbrace{P_\theta(T|S) + \mathbb{E}_{T' \sim P_\theta(T'|S)} [Q_\phi(S|T')] - Q_\phi(S|T)}_{\text{Adversarial Training}} \quad (5)$$

As we can see, we still maximize the mutual information when optimizing P_θ , but we minimize it when optimizing Q_ϕ . Moreover, we add the third term which minimizes the negative likelihood of Q_ϕ when feeding it with the real data. By regularizing the backward network with both the synthetic data and real data, the forward network will be trained against a smarter backward network that only rewards informative and related text. Here we omit the logarithm symbol for simplicity, and we also add $P_\theta(T|S)$ to the objective in order to balance the positive samples, which is referred as the teacher-forcing algorithm (Li et al., 2017).

3.1. Wasserstein Distance Minimization

Unfortunately, training such a framework similar to GAN is well known of being delicate and unstable (Arjovsky &

Bottou, 2017). The recent work of (Arjovsky et al., 2017) indicated the superiority of the *Wasserstein distance* in learning distributions due to its continuity and differentiability. They demonstrated their method significantly improves the stability of training and gained significant performance with the objective of minimizing the Wasserstein distance. In the following, we will prove that our adversarial training part of the objective function can be transformed to minimizing the Wasserstein distance under some simple assumptions.

To learn the proper distribution of the dialog utterances, in practice we replace the third term of the equation (5) by the expectation $\mathbb{E}_{T \sim P_r(D)} Q_\phi(S|T)$ where $P_r(D)$ is the real data distribution, D is the dataset.

Theorem 1. *Let $P_r(D)$ be any distribution over D . Let $P_\theta(\cdot|S)$ be the distribution of a neural network given S . $Q_\phi(S|T)$ represents the probability of a neural network where the input is T and the output is S , which can be formed as a function $f_\phi(S, T)$. We say the following problem can be considered as the Wasserstein Distance between P_r and P_θ if there is a constant K satisfying $\|f_\phi\|_L \leq K$:*

$$\min_{\phi} \mathbb{E}_{T' \sim P_\theta(T'|S)} [Q_\phi(S|T')] - \mathbb{E}_{T \sim P_r(D)} [Q_\phi(S|T)] \quad (6)$$

where $\|\cdot\|_L \leq K$ denotes K -Lipschitz.

Proof. It is obvious that the above problem is equivalent to:

$$\max_{\phi} \mathbb{E}_{T \sim P_r(D)} [Q_\phi(S|T)] - \mathbb{E}_{T' \sim P_\theta(T'|S)} [Q_\phi(S|T')] \quad (7)$$

which is also equivalent to solving the problem:

$$\sup_{\|f_\phi\|_L \leq K} \mathbb{E}_{T \sim P_r(D)} [Q_\phi(S|T)] - \mathbb{E}_{T' \sim P_\theta(T'|S)} [Q_\phi(S|T')] \quad (8)$$

where supremum is over all the K -Lipschitz functions Q_ϕ . According to the Kantorovich-Rubinstein duality (Villani, 2008), the above form is K times of the Wasserstein Distance between P_r and P_θ :

$$W(P_r, P_\theta) = \inf_{\gamma \in \Pi(P_r, P_\theta)} \mathbb{E}_{(x,y) \sim \gamma} [\|x - y\|] \quad (9)$$

where $\Pi(P_r, P_\theta)$ denotes the set of all joint distributions $\gamma(x, y)$ whose marginals are respectively P_r and P_θ . \square

In order to satisfy the K -Lipschitz striction of f_ϕ , we use the weight clipping trick (Arjovsky et al., 2017) to clamp the weights to a fixed box (say $[-0.01, 0.01^l]$) after each gradient update. Therefore, we modify the adversarial training part in the equation (5) as:

$$\min_{\theta} \max_{\phi} \mathbb{E}_{T \sim P_r(D)} [Q_\phi(S|T)] - \mathbb{E}_{T' \sim P_\theta(T'|S)} [Q_\phi(S|T')] \quad (10)$$

which can be transformed to the problem of minimizing a Wasserstein Distance.

We notice that sometimes the forward network may generate T' that is the same as T or very similar, which should not be punished for the adversarial training of the backward model. Hence, we add a multiplier $K(T')$ to the second term of the equation (10) to regularize the minimization procedure:

$$\begin{aligned} & - \mathbb{E}_{T' \sim P_\theta(T'|S)} [K(T') \cdot Q_\phi(S|T')] \\ & = - \mathbb{E}_{T' \sim P_\theta(T'|S)} [(1 - \cos(T, T')) \cdot Q_\phi(S|T')] \end{aligned} \quad (11)$$

where $\cos(T, T')$ is the cosine similarity between the embeddings of T and T' .

Since the equation (11) w.r.t. θ is not differentiable, we adopt Monte Carlo samples using the REINFORCE policy (Williams, 1992) to approximate the gradient with regard to θ :

$$\begin{aligned} & - \nabla_{\theta} \mathbb{E}_{T' \sim P_\theta(T'|S)} [K(T') \cdot Q_\phi(S|T')] \\ & = - \mathbb{E}_{T' \sim P_\theta(T'|S)} [K(T') \cdot Q_\phi(S|T') - b] \cdot \nabla_{\theta} \log P_\theta(T'|S), \end{aligned} \quad (12)$$

To remedy the high variance of learning signals, we adopt the baseline function b by empirically averaging the signals to stabilize the learning process (Williams, 1992). The gradient of the second term of the equation (11) with regard to ϕ can be calculated as:

$$\begin{aligned} & - \nabla_{\phi} \mathbb{E}_{T' \sim P_\theta(T'|S)} [K(T') \cdot Q_\phi(S|T')] \\ & = - \mathbb{E}_{T' \sim P_\theta(T'|S)} [K(T') \cdot \nabla_{\phi} Q_\phi(S|T')] \end{aligned} \quad (13)$$

3.2. Latent Noise Sampling

For the Monte Carlo sampling approaches in the policy gradient for NLP tasks, existing works often employ low-level sampling methods such as greedy search or beam search (Li et al., 2016b; Paulus et al., 2018). However, methods like beam search usually produce very similar or even identical results, thus require a large beam width and become inefficient and time-consuming.

Recent studies in the variational autoencoder (VAE) frameworks have demonstrated that sampling from high-level space can generate abundant variability of natural language by capturing its global and long-term structure (Serban et al., 2017; Zeng et al., 2019). Nevertheless, a known problem of the VAE-based dialog model is *degeneracy*, which ignores the latent variable and causes model to behave as a vanilla sequence-to-sequence model (Bowman et al., 2016; Park et al., 2018). Moreover, instead of modeling the relation between the given source and target pairs, we expect our samples to be diverse and not limited by the knowledge of the target label. Thus we sample the target text from $P_\theta(T'|S)$ by simply adding a noise to the latent representations (§2.1) of the encoder:

$$\tilde{z} = z + \delta. \quad (14)$$

Models	Relevance			Diversity		
	Average	Greedy	Extrema	Dist-1	Dist-2	Ent-4
HRED (Serban et al., 2016)	0.820	0.623	0.391	0.063	0.262	8.861
cGAN (Li et al., 2017)	0.841	0.648	0.410	0.069	0.272	9.418
VHCR (Park et al., 2018)	0.855	0.663	0.430	0.085	0.348	9.602
Dir-VHRED (Zeng et al., 2019)	0.862	0.660	0.431	0.076	0.306	9.586
LSTM	0.817	0.611	0.383	0.061	0.260	8.855
LSTM + MMI	0.825	0.617	0.391	0.065	0.265	8.894
LSTM + AMI (w/o LNS)	0.851	0.644	0.423	0.078	0.287	9.176
LSTM + AMI	0.858	0.655	0.429	0.082	0.294	9.205
Transformer	0.828	0.627	0.398	0.077	0.347	9.440
Transformer + MMI	0.835	0.636	0.406	0.082	0.356	9.493
Transformer + AMI (w/o LNS)	0.859	0.661	0.424	0.096	0.372	9.632
Transformer + AMI	0.869	0.672	0.433	0.104	0.385	9.695

Table 1. Quantitative evaluation for dialog generation on the PersonaChat dataset. “LNS” denotes the latent noise sampling. The top part presents the baselines, the medium and bottom parts show the performance and ablation results of our AMI framework based on the LSTM and Transformer. Our AMI significantly improves both LSTM and Transformer and achieves the state-of-the-art results.

Since the goal of the noise is to maximize both (1) the diversity the generated T' and (2) the probability of reconstructing S with the generated T' , we optimize it by:

$$\max_{\delta} \lambda \|\delta\| + \mathbb{E}_{T' \sim P_{\theta}(T'|S, \delta)} [Q_{\phi}(S|T')] \quad (15)$$

where λ is a scaling factor accounting for the difference in magnitude between the above two terms. The noise δ is computed by:

$$\delta \sim \mathcal{N}(\boldsymbol{\mu}(z), \boldsymbol{\sigma}^2(z)) \quad (16)$$

$$\boldsymbol{\mu}(z) = \mathbf{W}_2 \cdot \text{ReLU}(\mathbf{W}_1 z + \mathbf{b}_1) \quad (17)$$

$$\boldsymbol{\sigma}(z) = \text{Softplus}(\boldsymbol{\mu}(z)) \quad (18)$$

where $\mathcal{N}(\cdot, \cdot)$ is a Gaussian distribution. Since the backpropagation can not flow through a random node, in practice, we use the reparameterization trick (Kingma & Welling, 2014):

$$\delta = \boldsymbol{\mu}(z) + \epsilon \cdot \boldsymbol{\sigma}(z) \quad (19)$$

where the sample $\epsilon \sim \mathcal{N}(0, \mathbf{I})$. Therefore, the final AMI objective function is:

$$\min_{\theta} \max_{\phi} \mathbb{E}_{T \sim P_r(D)} [Q_{\phi}(S|T)] - \mathbb{E}_{T' \sim P_{\theta}(T'|S, \delta)} [K(T') \cdot Q_{\phi}(S|T')] - P_{\theta}(T|S) \quad (20)$$

where $T' \sim P_{\theta}(T'|S, \delta)$ means sampling the target text T' from $P_{\theta}(T'|S)$ by adding the noise δ in the high-level latent space of the model. Note that for the RNN-based sequence-to-sequence model, z denotes the concatenation of the last hidden states from both directions.

4. Experiments

To show the effectiveness of our approach, we use the bi-directional LSTM and Transformer (Vaswani et al., 2017) as our base architectures. We apply our AMI framework

to the dialog generation (§4.1) and neural machine translation (§4.2), which are representatives of two popular text generation tasks, and conduct comprehensive analyses on them. Code is available at <https://github.com/ZJULearning/AMI>.

4.1. Dialog Generation

We first verify the effectiveness of our method on the dialog generation task, which requires to generate a coherent and meaningful response given a conversation history.

Dataset We evaluate our dialog model on the PersonaChat dataset¹ (Zhang et al., 2018a). The dataset consists of conversations between crowdworkers who were randomly paired and asked to act the part of a given persona and chat naturally. There are around 160,000 utterances in around 11,000 dialogues, with 2000 dialogues for validation and test, which use non-overlapping personas.

Implementation Details We pre-train the forward network and backward network, and iteratively optimize the noise δ , forward network and backward network during AMI training. For the backward network, we follow (Li et al., 2016b) to predict the last utterance of the dialog history. Detailed configurations are in the Appendix B.1. To measure the relevance between the generated utterance and the ground-truth one, we use the three categories of word-embedding metrics (Rus & Lintean, 2012; Serban et al., 2017): *average*, *greedy* and *extrema*. The *average* metric calculates sentence-level mean embeddings, while the *greedy* and *extrema* compute the word-to-word cosine similarity. Their difference lies on that *greedy* takes the average word vector in the sentence as the sentence embedding while

¹<https://github.com/facebookresearch/ParlAI/tree/master/projects/personachat>

Backward Network	Synthetic	Real
MMI	0.807	0.805
MMI (fixed ϕ)	0.809	0.812
AMI	0.776	0.820

Table 2. Embedding average for the output of the backward network, where the input is the utterance (1) generated by the forward network or (2) from the dataset created by human.




Figure 2. Convergence curves for different ablated models. “regularizer” denotes the cosine multiplier in the equation (11).

extrema adopts the extreme value of these word vectors. To evaluate the diversity of the generated questions, we follow (Li et al., 2016a) to calculate $Dist-n$ ($n=1,2$), which is the proportion of unique n-grams over the total number of n-grams in the generated utterances for all conversations, and (Zhang et al., 2018b) to use the *Ent-4* metric, which reflects how evenly the n-gram distribution is over all generated utterances.

Baselines We compare our proposed method based on the BiLSTM and Transformer with the following baselines: (i) HRED (Serban et al., 2017): The hierarchical recurrent encoder-decoder baseline model. (ii) cGAN (Li et al., 2017): Conditional generative adversarial network based on the sequence-to-sequence model. (iii) VHCR (Park et al., 2018): Variational hierarchical conversation RNN, a hierarchical latent variable model with the utterance drop regularization. (iv) Dir-VHRED (Zeng et al., 2019): Using Dirichlet distribution in place of traditional Gaussian distribution in variational HRED for dialogue generation.

Results In Table 1, we apply our AMI framework to LSTM and Transformer, and compare them with other competitive baseline models on the PersonaChat dataset. As we can see, our method with Transformer clearly outperforms all the baselines and achieves state-of-the-art results. From the medium and bottom parts, we observe that our AMI greatly improves the relevance and diversity performances over both the LSTM and Transformer. Compared

Opponent Models	Wins	Losses	Ties
LSTM + AMI vs. LSTM + MMI	0.62	0.24	0.14
TF + AMI vs. TF + MMI	0.54	0.28	0.18
TF + AMI vs. Dir-VHRED	0.36	0.30	0.34
Dir-VHRED vs. Human	0.28	0.56	0.16
LSTM + AMI vs. Human	0.24	0.54	0.22
TF + AMI vs. Human	0.32	0.48	0.20

Table 3. Human evaluation results for the dialog generation. “TF” means the Transformer, “Human” means the original human-created dialog utterances in the dataset. Both “Wins” and “Losses” point to the left models.

to MMI, our method significantly exceeds on both the relevance and diversity metrics, which indicates that our adversarial training scheme helps improve the quality of the backward model, which subsequently provides more credible rewards to the forward network. We also conduct ablation experiments that when we use beam search for sampling instead of our proposed latent sampling noise, we find that the results drop on both metrics, which demonstrates that this sampling strategy can enlarge the searching space from the high-level structure while maximizing the mutual information between the output and the source text.

As shown in Table 2, we compare the performance of the backward network before and after training by our AMI framework. We evaluate the embedding average of the outputs of the backward network when feeding it with the synthetic data and the real data. Compared to the MMI with fixed ϕ (i.e. the pre-trained model), the performance of the regular MMI backward model drops and its output text is more relevant to ground truth when fed with the synthetic data rather than the real data. This means such a training framework misleads the backward model to give lower rewards to the generated targets that look more like the real ones, and thus results in negative effect to the optimization of the forward model. With our adversarial training scheme, the backward model performs better when fed with the real data and worse with the synthetic data, which means it improves the ability to manage the information flow.

We also present the convergence curves over different ablated models of the AMI for LSTM in Figure 2. We can see that the regularizer $K(T')$ (equation (11)) helps accelerate the convergence of the model. We also observe that without the weight clipping, which is important in transforming our objective to minimizing the Wasserstein distance, the training becomes more unstable and the loss can not approach the performance of the full AMI.

Human Evaluation We conduct the human evaluation to measure the quality of the generated dialog utterances, as shown in Table 3. Models are paired as opponents and we randomly presented 120 dialog history and the outputs of




Figure 3. Variational lower bound over training steps.

each paired models to 5 judges, who are asked to decide which of the two outputs is better. Ties are permitted. From the human evaluated results, our AMI is proven to be better at generating a coherent and meaningful dialog response than MMI, and our AMI on Transformer performs better than Dir-VHRED, one of the state-of-the-art dialog models. However, all of the synthetic utterances lost in the comparisons with the human created sentences, which indicates that the current generative models still have much room for improvement in understanding and generating natural language.

Bound Analysis To evaluate whether the mutual information between the source and target text can be efficiently maximized, or whether our method can lead to a tighter variational lower bound, we compare the variational information maximization (VIM) lower bounds over MMI and our proposed AMI framework, both in Transformer architectures. We see from Figure 3 that the lower bound in AMI is higher than in MMI. According to the equation (3), if the backward models Q_ϕ of AMI and MMI get the same KL-divergence value, then the AMI obtains a higher mutual information between the source and target than MMI; On the other hand, if AMI and MMI get the same mutual information, the AMI obtains a lower KL-divergence of the backward model than the MMI, which means the backward network in the AMI is more closed to the posterior distribution, thus the variational lower bound becomes tighter.

Case Study In Table 4, we show two examples of the generated dialog responses along with the dialog history. In the first case, the MMI model tries to improve communication by adding a question. However, the added question “what is your favorite color” has nearly no connection with the previous dialog history. Our method provides a more informative answer to the question of the role [B], and raises a very reasonable question in this situation. In the second case, the Transformer baseline produces a very common question, which can just carry the conversation but is not

[A]: Oh , we had a graduation party from college for me.

[B]: Congratulations ! What is your degree in?

Transformer: I am studying to become a fashion designer.

Transformer + MMI: I am studying law. What is your favorite color?

Transformer + AMI: I am majoring in accounting and minoring in business management. What about you?

Human: Math. Also had a wedding reception for my brother.

[A]: What is your novel about? I am a recent college graduate.

[B]: It is actually a romance novel, but it has sci-fi elements.

Transformer: What do you do for a living?

Transformer + MMI: I do not really have a favorite. I do not mind.

Transformer + AMI: Oh, I see. I am a huge fan of romance books.

Human: That sounds interesting, I would like to read it some day.

Table 4. Examples of dialog generation. “Human” denotes the original response in the dataset, and all the generation models represent the role [A].

coherent and has little mutual information with the previous dialog history. The MMI provides a response that has more connection with the dialog history but is less suitable for this conversation since no one asked about his/her preference. Our method, however, generates a response that is not only informative but also coherent to the previous dialog.

4.2. Neural Machine Translation

We also evaluate our method on the neural machine translation (NMT) task, which requires to encode a source language sentence and predict a target language sentence.

Dataset The dataset for our evaluation is the WMT translation task between English and German in both directions and the translation performances are reported on the official test set `newstest2014`². The dataset contains 4.5M training pairs, 2169 validation pairs and 3003 test pairs.

Implementation Details The training scheme for each module is the same as the dialog generation task. We use BLEU (Papineni et al., 2002) as the evaluation metric for translation quality. BLEU is the geometric mean of n -gram precisions where $n \in \{1, 2, 3, 4\}$ weighted by sentence lengths. Detailed configurations are in the Appendix B.2.

Results We present the performance of our AMI for machine translation in Table 5. Similar to the results of dialog

²<http://www.statmt.org/wmt14/translation-task.html>

Models	En→De	De→En
LSTM	23.1	27.0
LSTM + MMI	23.3	27.3
LSTM + AMI (w/o LNS)	24.8	28.9
LSTM + AMI	25.0	29.1
Transformer	26.7	30.2
Transformer + MMI	26.9	30.5
Transformer + AMI (w/o LNS)	28.2	31.9
Transformer + AMI	28.5	32.1

Table 5. Machine translation results of the BLEU scores on WMT English-German for newstest2014.

Ein weiteres Radarsensor prüft, ob die Grünphase für den Fußgänger beendet werden kann.

Transformer: Another radar sensor checks whether the green phase can be terminated for pedestrians.

Transformer + MMI: Another radar sensor checks whether the green phase can be closed for the pedestrian.

Transformer + AMI: Another radar sensor checks whether the green phase for pedestrians can be stopped.

Human: An additional radar sensor checks whether the green phase for the pedestrian can be ended.

Table 6. An example of German-to-English machine translation. “Human” denotes the original sentence in the dataset.

generation, our AMI framework improves upon both the LSTM and Transformer, and outperforms the MMI for about 1.4 BLEU scores. We also find that the improvement from LNS is quite modest, and we conjecture this is because the NMT task has the nature that the semantics of the target are fully specified by the source.

Case Studies In Table 6, we present an example of the German-to-English translation over different models. As we can see, all of the three models correctly translate almost all the words. However, both the Transformer and Transformer with MMI translate the goal of “ending the green phase” to “for pedestrian”, which makes the meanings of the generated sentences totally different from the original one. However, our method produces almost the same sentence with the ground truth one except for a few synonyms.

5. Related Works

Estimating mutual information (Bahl et al., 1986; Brown, 1987) has been comprehensively studied for many tasks such as Bayesian optimal experimental design (Ryan et al., 2016; Foster et al., 2018), image caption retrieval (Mao et al., 2015), neural networks explanation (Tishby et al., 2000; Tishby & Zaslavsky, 2015; Gabrié et al., 2018), etc. However, adapting MMI to sequence modeling such as text generation is empirically nontrivial as we typically have

access to discrete samples but not the underlying distributions (Poole et al., 2019). Li et al. (2016a) proposed to use MMI as the objective function to address the issue of output diversity in the neural generation framework. However, they use the MI-prompting objective only for testing, while the training procedure remains the same as the standard MLE. Li et al. (2016b) addressed this problem by using deep reinforcement learning to set the mutual information as the future reward. Zhang et al. (2018b) learned a dual objective to simultaneously learn two mutual information between the forward and backward models. Ye et al. (2019) proposed the dual information maximization to jointly model the dual information of two tasks. However, they optimize the backward model in the same direction with the forward model, which would limit its approach to the true posterior distribution, thus result in an unreliable reward for the forward model.

Recent development of dual learning leverages the similar idea of enforcing forward-backward consistency in language translation (He et al., 2016; Artetxe et al., 2018) and image-to-image translation (Zhu et al., 2017; Yi et al., 2017). They minimize the reconstruction loss of backward translation to verify and improve translation quality. Unlike these works, our backward network is trained as an auxiliary model, which aims to play against the forward network to iteratively promote generation of more meaningful text.

Adversarial training for text generation has also gained significant popularity (Yu et al., 2017; Li et al., 2017; Zhang et al., 2017; Che et al., 2017; Lin et al., 2017; Guo et al., 2018; Chen et al., 2018). The idea behind all these works is to use the real and synthetic data to train a classifier as the discriminator, which challenges the generator to produce text that looks more natural. However, some safe but uninformative phrases like “I don’t know” or “can you say it again” can be judged as human-generated in many instances especially when they appear frequently in the dataset, which would be easier to learn by the discriminator. Therefore, in this paper, we use the mutual information as the reward, which is much stricter than a binary classifier because it will give positive signals only if the backward model can reconstruct the source text from the generated target text.

6. Conclusion

In this paper, we introduced *Adversarial Mutual Information* (AMI), a novel text generation framework that addresses a minimax game to iteratively learn and optimize the mutual information between the source and target text. The forward network in the framework is trained by playing against the backward network that aims to reconstruct the source text only if its input is in the real target distribution. We proved that our objective can be trained more stably by being transformed to the problem of Wasserstein distance

minimization. The experimental results on two popular text generative tasks demonstrated the effectiveness of our framework, and we show our method has the potential to lead a tighter lower bound of the MMI problem. In future, we will attempt to explore a lower variance and more unbiased gradient estimator for the text generator in this framework and apply the AMI in multi-modality situations.

Acknowledgements

This work was supported in part by The National Key Research and Development Program of China (Grant Nos: 2018AAA0101400), in part by The National Nature Science Foundation of China (Grant Nos: 61936006), in part by the Alibaba-Zhejiang University Joint Institute of Frontier Technologies, in part by the China Scholarship Council, in part by the AWS Machine Learning Research Awards, and in part by U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

References

- Akoury, N., Krishna, K., and Iyyer, M. Syntactically supervised transformers for faster neural machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1269–1281, 2019.
- Arjovsky, M. and Bottou, L. Towards principled methods for training generative adversarial networks. In *5th International Conference on Learning Representations*, 2017.
- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein generative adversarial networks. In *International conference on machine learning*, pp. 214–223, 2017.
- Artetxe, M., Labaka, G., Agirre, E., and Cho, K. Unsupervised neural machine translation. In *International Conference on Learning Representations*, 2018.
- Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- Bahl, L., Brown, P., De Souza, P., and Mercer, R. Maximum mutual information estimation of hidden markov model parameters for speech recognition. In *ICASSP’86. IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 11, pp. 49–52. IEEE, 1986.
- Barber, D. and Agakov, F. V. The im algorithm: a variational approach to information maximization. In *Advances in neural information processing systems*, pp. None, 2003.
- Bowman, S. R., Vilnis, L., Vinyals, O., Dai, A., Jozefowicz, R., and Bengio, S. Generating sentences from a continuous space. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pp. 10–21, 2016.
- Brown, P. F. The acoustic-modeling problem in automatic speech recognition. Technical report, CARNEGIE-MELLON UNIV PITTSBURGH PA DEPT OF COMPUTER SCIENCE, 1987.
- Che, T., Li, Y., Zhang, R., Hjelm, R. D., Li, W., Song, Y., and Bengio, Y. Maximum-likelihood augmented discrete generative adversarial networks. *arXiv preprint arXiv:1702.07983*, 2017.
- Chen, L., Dai, S., Tao, C., Zhang, H., Gan, Z., Shen, D., Zhang, Y., Wang, G., Zhang, R., and Carin, L. Adversarial text generation via feature-mover’s distance. In *Advances in Neural Information Processing Systems*, pp. 4666–4677, 2018.
- Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., and Abbeel, P. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems*, pp. 2172–2180, 2016.
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pp. 1724–1734, 2014.
- Foster, A., Jankowiak, M., Bingham, E., Teh, Y. W., Rainforth, T., and Goodman, N. Variational optimal experiment design: Efficient automation of adaptive experiments. *NeurIPS Bayesian Deep Learning Workshop*, 2018.
- Gabrié, M., Manoel, A., Luneau, C., Macris, N., Krzakala, F., Zdeborová, L., et al. Entropy and mutual information in models of deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 1821–1831, 2018.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.
- Guo, J., Lu, S., Cai, H., Zhang, W., Yu, Y., and Wang, J. Long text generation via adversarial training with leaked information. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- He, D., Xia, Y., Qin, T., Wang, L., Yu, N., Liu, T.-Y., and Ma, W.-Y. Dual learning for machine translation. In *Advances in neural information processing systems*, pp. 820–828, 2016.

- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *ICLR*, 2014.
- Klein, G., Kim, Y., Deng, Y., Senellart, J., and Rush, A. M. Opennmt: Open-source toolkit for neural machine translation. *arXiv preprint arXiv:1701.02810*, 2017.
- Li, J., Galley, M., Brockett, C., Gao, J., and Dolan, B. A diversity-promoting objective function for neural conversation models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 110–119, 2016a.
- Li, J., Monroe, W., Ritter, A., Jurafsky, D., Galley, M., and Gao, J. Deep reinforcement learning for dialogue generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1192–1202, 2016b.
- Li, J., Monroe, W., Shi, T., Jean, S., Ritter, A., and Jurafsky, D. Adversarial learning for neural dialogue generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 2157–2169, 2017.
- Lin, K., Li, D., He, X., Zhang, Z., and Sun, M.-T. Adversarial ranking for language generation. In *Advances in Neural Information Processing Systems*, pp. 3155–3165, 2017.
- Mao, J., Xu, W., Yang, Y., Wang, J., Huang, Z., and Yuille, A. Deep captioning with multimodal recurrent neural networks (m-rnn). *ICLR*, 2015.
- Pan, B., Yang, Y., Li, H., Zhao, Z., Zhuang, Y., Cai, D., and He, X. Macnet: Transferring knowledge from machine comprehension to sequence-to-sequence models. In *Advances in Neural Information Processing Systems*, pp. 6092–6102, 2018.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pp. 311–318. Association for Computational Linguistics, 2002.
- Park, Y., Cho, J., and Kim, G. A hierarchical latent structure for variational conversation modeling. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 1792–1801, 2018.
- Paulus, R., Xiong, C., and Socher, R. A deep reinforced model for abstractive summarization. In *International Conference on Learning Representations*, 2018.
- Poole, B., Ozair, S., Van Den Oord, A., Alemi, A., and Tucker, G. On variational bounds of mutual information. In *International Conference on Machine Learning*, pp. 5171–5180, 2019.
- Rus, V. and Lintean, M. A comparison of greedy and optimal assessment of natural language student input using word-to-word similarity metrics. In *Proceedings of the Seventh Workshop on Building Educational Applications Using NLP*, pp. 157–162, 2012.
- Ryan, E. G., Drovandi, C. C., McGree, J. M., and Pettitt, A. N. A review of modern computational algorithms for bayesian optimal design. *International Statistical Review*, 84(1):128–154, 2016.
- See, A., Liu, P. J., and Manning, C. D. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pp. 1073–1083, 2017.
- Sennrich, R., Haddow, B., and Birch, A. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1715–1725, 2016.
- Serban, I. V., Sordoni, A., Bengio, Y., Courville, A., and Pineau, J. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- Serban, I. V., Sordoni, A., Lowe, R., Charlin, L., Pineau, J., Courville, A., and Bengio, Y. A hierarchical latent variable encoder-decoder model for generating dialogues. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- Srivastava, N., Hinton, G. E., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- Sutskever, I., Vinyals, O., and Le, Q. V. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pp. 3104–3112, 2014.
- Tishby, N. and Zaslavsky, N. Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop (ITW)*, pp. 1–5. IEEE, 2015.
- Tishby, N., Pereira, F. C., and Bialek, W. The information bottleneck method. *arXiv preprint physics/0004057*, 2000.

- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Villani, C. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3-4):229–256, 1992.
- Wu, C.-S., Socher, R., and Xiong, C. Global-to-local memory pointer networks for task-oriented dialogue. In *International Conference on Learning Representations*, 2019.
- Ye, H., Li, W., and Wang, L. Jointly learning semantic parser and natural language generator via dual information maximization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 2090–2101, 2019.
- Yi, Z., Zhang, H., Tan, P., and Gong, M. Dualgan: Un-supervised dual learning for image-to-image translation. In *Proceedings of the IEEE international conference on computer vision*, pp. 2849–2857, 2017.
- Yu, L., Zhang, W., Wang, J., and Yu, Y. Seqgan: Sequence generative adversarial nets with policy gradient. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- Zeng, M., Wang, Y., and Luo, Y. Dirichlet latent variable hierarchical recurrent encoder-decoder in dialogue generation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*, pp. 1267–1272, 2019.
- Zhang, S., Dinan, E., Urbanek, J., Szlam, A., Kiela, D., and Weston, J. Personalizing dialogue agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, pp. 2204–2213, 2018a.
- Zhang, Y., Gan, Z., Fan, K., Chen, Z., Henao, R., Shen, D., and Carin, L. Adversarial feature matching for text generation. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 4006–4015, 2017.
- Zhang, Y., Galley, M., Gao, J., Gan, Z., Li, X., Brockett, C., and Dolan, B. Generating informative and diverse conversational responses via adversarial information maximization. In *Advances in Neural Information Processing Systems*, pp. 1815–1825, 2018b.
- Zhu, J.-Y., Park, T., Isola, P., and Efros, A. A. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pp. 2223–2232, 2017.